Bandolier 95

What do we think?
What do we know?
What can we prove?

Evidence-based health care **£3.00** **January 2002 Volume 9 Issue 1**

Choosing items for ***Bandolier*** is a bit like a random walk, peppered with signposts to interesting places, laden down with guidebooks, and with folk continually coming up to tell you about what you've already missed in your travels, or with exciting possibilities about where you might go next. But it's also a random walk with a history, and ***Bandolier***'s history is now 94 issues with a massive Internet site. Despite that, there's always *something* that connects items or issues and magically pulls them together.

Deciding by talking

In this first issue of 2002 that *something* was discussions with GPs in November and December. How many cases of Perthes' disease should we expect in our PCO every year? How does our PCO ensure that we get value for money from prescribing? How important is compliance? Why do prescribing advisors sometimes insist that we GPs do things we know to be crazy? How do we know whether we are treating patients with different conditions equitably?

Enough here for a lifetime, let alone an issue. There's masses written on each of these, of course, but modern electronic wizadry, and a bit of luck, enables all sorts of goodies to be pulled from the ether to help.

Brains in gear for new year

For instance, the equity question is helped by some exceptional work from Scotland following up for five years everyone who in 1991 had an initial hospital admission for common cardiac conditions or cancer. Policy making about prescribing will rarely be answered in one concise little paper, but ***Bandolier*** found several that help. From Finlay McAlister, David Sackett and co in Canada come some typically thoughtful pieces on how we might think about class effects, and what we should have in mind when faced with equivalence trials. Both of these relate directly to prescribing policies.

Then there's a superb randomised trial and commentary about using different SSRIs for depression in general practice, and finally some British health economists highlight important thinking about the cost-effectiveness of trying to change professional behaviour. Powerful stuff when combined. Oh, and there's terrific information from Northern Ireland and Liverpool that helps to tell us how many cases of Perthes' disease to expect in a PCO.

Responding to questions asked by professionals is the most rewarding thing ***Bandolier*** does. Please keep them coming in 2002.

SURVIVAL WITH COMMON CANCERS AND HEART CONDITIONS

Scotland has produced many excellent people and ideas. Few, if asked, would number among them the Scottish Morbidity Record Scheme. In this scheme, each person admitted to a Scottish hospital since 1981 has been assigned a unique identifying number, used for all subsequent hospital admissions. These are also linked to diagnostic codes, and to the General Register office relating to all deaths in the UK, and the Scottish Cancer Registry. This results in a fantastic tool for examining disease and outcomes. It has been used to examine the five-year survival after initial hospital admission for common cardiac conditions and cancers [1].

Methods

Information for all of Scotland (population about 5 million) for first hospital admissions for 1991 were used. Excluded were patients who had an admission for their index condition in the preceding 10 years, and, for cancer patients, any with an admission for any malignant neoplasm.

All deaths occurring in individuals before their expected age of death, taken from life-expectance tables for age-matched populations in 1991, were defined as premature. The number of expected life-years lost was calculated by subtracting the actual age at death from the expected age of death. Loss of expected life was then calculated as a median for each diagnosis and for a 1000 population.

Results

There were 14,842 initial diagnoses for women of heart failure, myocardial infarction, and breast, lung, bowel and ovarian cancer. There were 16,224 initial diagnoses for men for heart failure, myocardial infarction, and lung, bowel,

In this issue

Survival after first hospital admission ..... p. 1
Class and equivalence ..... p. 3
Equal on average is not equal for everyone ..... p. 6
Cost of change - why bother? ..... p. 7
Perthes' disease incidence ..... p. 7

www.ebandolier.com

Table 1: After first hospital admission with heart condition or common cancer in Scotland, 1991. Incidence, five year survival, and expected life years lost per person and per 1000 population

Sex/condition	Number	Mean age	Annual incidence per 1000	Five year survival (%	Median expected life-years lost per person	Expected life years lost per 1000
Women						
Heart failure	3606	71	1.4	25	6.8	5.1
Myocardial infarction	4916	64	1.9	48	7.9	6.7
Lung cancer	2902	69	0.8	5	13.1	6.7
Breast cancer	1490	69	0.4	65	16.5	7.0
Large bowel cancer	1402	74	0.4	35	10.2	3.0
Ovarian cancer	526	69	0.2	30	14.6	2.3
Men						
Heart failure	3241	76	1.3	25	8.7	6.8
Myocardial infarction	6932	72	2.8	60	9.7	9.4
Lung cancer	2695	62	0.8	5	14.4	12.3
Large bowel cancer	1385	70	0.6	35	10.3	3.6
Prostate cancer	1211	72	0.5	38	5.6	1.2
Bladder cancer	760	64	0.3	52	6.7	0.9

prostate and bladder cancer. The numbers, mean age and annual incidence are shown in Table 1.

Also shown are the five-year survivals (read from graphs, so limited accuracy here), and the median expected life years lost to individual patients with a diagnosis and for a population of 1000 (adjusted for the proportion of deaths that were premature).

Common cardiac conditions in men and women were more common than cancers, but were associated with similar five-year survival rates and life years lost. Five-year survival for heart failure in men and women was associated with a low survival rate (25%), even when adjusted for age. The age-specific probability of surviving five years for the population was about 80% for women and 75% for men.

Comment

The fact that common cardiac conditions and cancers are associated with poor outcomes is hardly a matter of surprise. What is interesting here is that we have reasonable numbers in a whole population and with information collected systematically. We can compare societal impacts of the various diseases on mortality directly.

The heart failure angle

There is more than just interesting comparative figures in this paper, though. It also contains an argument why we should probably be doing more about heart failure, and doing more more effectively. In doing so it reviews, for instance, other studies in other developed countries, and demonstrates that survival rates in Scotland were broadly similar to those in the USA, Australia and other European countries. This is not just a Scottish problem.

It goes on to review briefly the burden of the disease, and the benefits of screening and palliative care programs, making comparisons with cancer screening and screening for malignant hypertension. It finally reviews specific initiatives that could be used in heart failure programs.

There is much food for thought here. The authors' interest is in heart failure, and their discussion is worth reading because it is wide ranging and intelligent. They make a point that better results can be achieved with heart failure, and subsequent hospital admissions avoided, by the use of nurse-led, comprehensive management programmes. But those, in PCOs or elsewhere, who have to plan and organise services will find this rewarding and useful background reading. More than anything else, this is one of those papers that makes you sit back and have a quiet think about what you are doing and why.

References:

1 S Stewart et al. More 'malignant' than cancer? Five-year survival following a first admission for heart failure. European Journal of Heart Failure 2001 2: 315-322.

CLASS AND EQUIVALENCE

Class (noun); any set of people or things grouped together or differentiated from others. An increasingly asked question is that of whether a set of drugs forms a class, and whether there is a 'class effect'. Class effect is usually taken to mean similar therapeutic effects and similar adverse effects, both in nature and extent. If such a 'class effect' exists, then it makes decision-making easy: you choose the cheapest.

Criteria for drugs to be grouped together as a class involve some or all of the following:

- Drugs with similar chemical structure
- Drugs with similar mechanism of action
- Drugs with similar pharmacological effects

Declaring a class effect requires a bit of thought, though. How much thought, and of what type, has been considered in one of that brilliant JAMA series on users guides to the medical literature [1]. No one should declare a class effect and choose the cheapest without reference to the rules of evidence set out in this paper.

Levels of evidence for efficacy

These are shown in Table 1, though if it comes down to levels 3 and 4 evidence for efficacy, the ground is pretty shaky. Level 1 evidence is what we always want and almost always never get, the large randomised head to head comparison. By the time there are enough compounds around to form a class, there is almost no organisation interested in funding expensive, new, trials to test whether A is truly better than B.

Most of the time we will be dealing with randomised trials of A versus placebo or standard treatment and B versus placebo or standard treatment. This will be level 2 evidence based on clinically important outcomes (a healing event) or validated surrogate outcomes (reduction of cholesterol with a statin). So establishing a class effect will likely involve quality systematic review or meta-analysis of quality randomised trials.

What constitutes quality in general is captured in Table 1, though there will be some situation-dependent factors. One thing missing from Table 1 is size. There probably needs to be some prior estimate of how many patients or events constitutes a reasonable number for analysis.

Levels of evidence for safety

These are shown in Table 2. There are always going to be problems concerning rare, but serious, adverse events. The inverse rule of three tells us that if we have seen no serious adverse events in 1500 exposed patients, then we can be 95% sure that they do not occur more frequently than 1 in 500 patients.

Randomised trials of efficacy will usually be underpowered to detect rate, serious adverse events, and we will usu-

Table 1: Levels of evidence for efficacy for class effect

Level	Comparison	Patients	Outcomes	Criteria for validity
1	RCT direct comparison	Identical	Clinically important	Randomisation concealment Complete follow up Double-blinding Outcome assessment must be sound
2	RCT direct comparison	Identical	Valid surrogate	Level 1 plus Validity of surrogate outcome
2	Indirect comparison with placebo from RCTs	Similar or different in disease severity or risk	Clinically important or valid surrogate	Level 1 plus Differences in methodological quality End points Compliance Baseline risk
3	Subgroup analyses from indirect comparisons of RCTs with placebo	Similar or different in disease severity or risk	Clinically important or valid surrogate	Level 1 plus Multiple comparisons, post hoc data dredging Underpowered subgroups Misclassification into subgroups
3	Indirect comparison with placebo from RCTs	Similar or different in disease severity or risk	Unvalidated surrogate	Surrogate outcomes may not capture all good or bad effects of treatment
4	Indirect comparison of nonrandomised studies	Similar or different in disease severity or risk	Clinically important	Confounding by indication, compliance, or time Unknown or unmeasured confounders Measurement error Limited database, or coding systems not suitable for research

Table 2: Levels of evidence for safety for class effect

Level	Type of study	Advantages	Criteria for validity
1	RCT	Only design that permits detection of adverse effects when the adverse effect is similar to the event the treatment is trying to prevent	Underpowered for detecting adverse events unless specifically designed to do so
2	Cohort	Prospective data collection, defined cohort	Critically depends on follow up, classification and measurement accuracy
3	Case-control	Cheap and usually fast to perform	Selection and recall bias may provide problems, and temporal relationships may not be clear.
4	Phase 4 studies	Can detect rare but serious adverse events if large enough	No control or unmatched control Critically depends on follow up, classification and measurement accuracy
5	Case series	Cheap and usually fast	Often small sample size, selection bias may be a problem, no control group
6	Case report(s)	Cheap and usually fast	Often small sample size, selection bias may be a problem, no control group

ally have to use other study designs. In practice the difficulty will be that soon after new treatments are introduced there will be a paucity of data for these other types of study. Only rarely will randomised trials powered to detect rare adverse events be conducted.

Most new treatments are introduced after being tested on perhaps a few thousand patients in controlled trials. Caution is needed in treatments for chronic conditions, especially difficult if trials are only short-term and where other diseases and treatments are likely.

Compliance

A difficult issue this, with a fragmented literature. But we do know that while compliance is usually high in clinical trials it may be lower in practice. Treatment schedules that are likely to improve compliance (once a day, for instance) might be important.

Cost

Economic studies are complicated beasts, and we need to treat this evidence with caution. Assumption of a class effect is usually done to justify choosing the cheapest drug in terms of acquisition (prescribing) costs. Terrific if this means that the costs of achieving the same ends are minimised. It may not be like that, and health economics in class effects need to be carefully thought through.

Comment

This paper uses statins as an example, with a decision being taken by clinician and policymaker between older, more expensive statins, and newer, cheaper, statins. Tactfully one chooses the cheaper statin with less information, and the other the older and more expensive statin with masses of patient experience. Can you guess who chose what?

Bandolier 47 examined the evidence for some of the older statins, with up to 27,000 years of patient experience and made the point that weight of evidence should be as important as acquisition cost. Having this paper to hand at the time would have been a great help.

Equivalence

McAlister & Sackett extend their thoughts on class effects to the particular example of equivalence trials, and provide some useful guides about what features of equivalence trials are important in determining their validity [2]. The intellectual problem with equivalence (A versus B) trials is that the same result is consistent with three conclusions:

1 Both A and B are equally effective
2 Both A and B are equally ineffective
3 Trials inadequate to detect differences between A and B

To combat the problems posed by the latter two conclusions, McAlister & Sackett suggest several criteria in addition to those used for superiority trials (A and / or B versus placebo). These are shown in Table 3.

Control shown previously to be effective?

Ideally documented in a systematic review of placebo controlled trials with benefits on active drug exceeding a clinically important effect. Without this information both may be equally ineffective.

Patients and outcomes similar to original trials?

Obvious, this one. If they are not, then any conclusion about equivalence is doomed. Beware, though, trials designed to show equivalent efficacy being used to demonstrate differences in harm or toxicity, for which they were not powered.

Table 3: Evidence quality for superiority and active-control equivalence trials

Superiority trials	Active-control equivalence trials
Randomised allocation	Randomised allocation
Randomisation concealed	Randomisation concealed
All patients randomised accounted for	All patients randomised accounted for
Intention to treat analysis	Intention to treat analysis **and on-treatment analysis**
Clinicians and patients blinded to treatment received	Clinicians and patients blinded to treatment received
Groups treated equally	Groups treated equally
Groups identical at baseline	Groups identical at baseline
Clinically important outcomes	Clinically important outcomes
	Active control previously shown to be effective
	Patients and outcomes similar to trials previously showing efficacy
	Both regimens applied in an optimal fashion
	Appropriate null hypothesis tested
	Equivalence margin pre-specified
Trial of sufficient size	Trial of sufficient size

Regimens applied in identical fashion?

The most common example is that of choosing the best dose of A versus an ineffective dose of B (no names, no pack drill, but no prizes for picking out numerous examples especially from pharmaceutical company sponsored trials showing "our drug is better than yours"). Should be OK if licensed doses are chosen.

Other pitfalls to look out for are low compliance or frequent treatment changes, incomplete follow up, disproportionate use of cointerventions and lack of blinding.

Appropriate statistical analysis?

Equivalence trials are designed to rule out meaningful differences between two treatments. Often one-sided tests of difference are used. Lack of significant superiority is not necessarily the same as defining an appropriate level of equivalence and testing for it.

Intention to treat analysis confers the risk of making a false-negative conclusion that treatments have the same efficacy when they do not. In equivalence trials the conservative approach may be to compare patients actually on treatment. Both analyses should probably be used.

Prespecified equivalence margin?

How different is different? Equivalence trials should have a prior definition of how big a difference is a difference, and justify it. Even more than that, they have to convince you that the lack of that difference means that treatments would, in fact, be equivalent.

Size?

Most equivalence trials do not have enough power to detect even a 50% difference between treatments, and a 1994 review [3] found that 84% were too small to detect a 25% difference. Size is everything when we want to show no difference, and the smaller the difference that is important, the larger the trial has to be.

Comment

McAlister & Sackett apply their methodological criteria to four large equivalence trials in hypertension. All had failings, and none could detect a 10% difference between treatments. Readers of equivalence trials should beware.

Designating a class effect on a group of drugs, and judging them to be equivalent on inadequate evidence is something most of us do at some time or another. Because prescribing costs often drive decisions, "cheapest is best" thinking often applies. Much of the time we will make incorrect decisions, but fortunately won't have the evidence to know that we are wrong. This is important and tricky territory that needs more work.

References:

1 FA McAlister et al. Users' guides to the medical literature XIX Applying clinical trial results B. Guidelines for determining whether a drug is exerting (more than) a class effect. JAMA 1999 282: 1371-1377.

2 FA McAlister & DL Sackett. Active-control equivalence trials and antihypertensive agents. American Journal of Medicine 2001 111: 553-558.

3 D Moher et al. Statistical power, sample size, and their reporting in randomized controlled trials. JAMA 1994 272: 122-124.

EQUAL ON AVERAGE DOES NOT MEAN EQUAL FOR EVERYONE

This fascinating title comes from a thoughtful commentary [1] on an even more thoughtful and useful trial [2] examining the effectiveness of SSRIs in a naturalistic trial in primary care. Together they make a useful contribution to thinking about effects of classes of drugs, of designs of trials to determine differences in practice, and how to use formularies.

Trial

The trial [2] was conducted in two primary care practice networks in the USA. The networks were of primary care physicians interested in optimising the care they provide through education and practice-based research.

The trial was designed to resemble real world practice. The decision to start an antidepressant was based strictly on physician judgement that there was clinical depression, rather than specific criteria for diagnosis. Patients and physicians knew what was being taken because blinding would obscure typical patient management. All decisions about discontinuation, switching, or dose change were made by physician and patient, with no pre-set criteria. There was no interaction between information gathering and physicians, who were uninfluenced by any results.

Apart from very obvious exclusions, adult patients with clinical depression were randomised (with adequate concealment) to initial doses of paroxetine (20 mg), fluoxetine (20 mg) or sertraline (50 mg). Drugs were free to patients. After enrolment, they were interviewed by telephone and questionnaire at several times over nine months. The primary outcome was SF-36 mental component summary, but a battery of other measures was used. Compliance (switching, discontinuation) was assessed.

Results

There were no differences between the three groups at baseline. Patients were predominantly white women with major depression, with an average age of 46 years.

All three SSRIs showed substantial improvements in depression and other outcomes. There were no differences between treatments for any outcome, using any definition of intention to treat or per protocol analysis, or with any subgroup analysis. In the whole sample, the proportion meeting criteria for major depression fell from 74% at baseline to 26% at nine months (Figure 1). Rapid initial benefit was followed by slow continued improvement.

Figure 1: Major depression with treatment

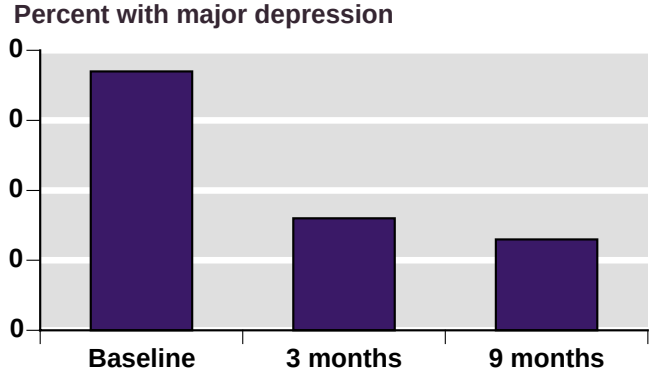

Figure 2: Stopping or switching

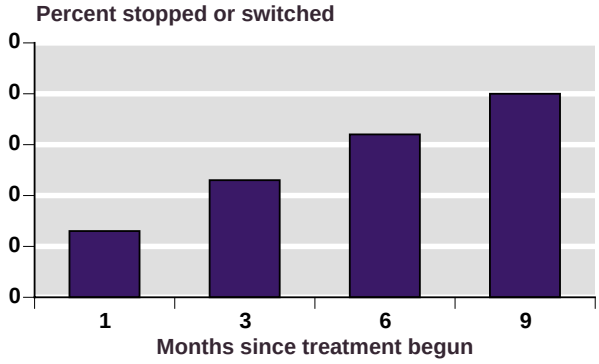

Patients initially randomised to one treatment frequently changed treatment. By nine months only 44% were still taking the treatment to which they had been randomised. Some (about 15%) were lost to follow up after baseline or when on randomised treatment. Others either switched to another antidepressant or stopped treatment because of adverse effects or lack of efficacy, again without any difference between the three SSRIs (Figure 2).

Comment

The study was not a formal equivalence study, but rather one that sought to mimic clinical practice, and was powered to detect a moderate change in outcome. It found no practical difference between the three SSRIs. It also found that clinical practice, in which dose or drug changes were done on the basis of patient-physician interaction as they would be in real life, resulted in many patients stopping or switching drugs.

As the accompanying commentary says, all other considerations being equal, an initial choice cased on prescription costs is prudent, ethical, and clinically reasonable. The study is important because it provides evidence for formulary choices for SSRIs in this indication, and indicates how evidence for other choices could be determined.

However, it also provides evidence against formularies that restrict use of SSRIs for subsequent choices when an initial choice is unsatisfactory. Because these three SSRIs are equal on average does not mean that they equal for every individual. There may be a class effect, but class effects are average effects, and the thinking here supports taking individuals into account.

References:

1 G Simon. Choosing a first line antidepressant. Equal on average does not mean equal for everyone. JAMA 2001 286: 3003-3004.

2 K Kroenke et al. Similar effectiveness of paroxetine, fluoxetine and sertraline in primary care. JAMA 2001 286: 2947-2995.

COST OF CHANGE – OR WHY BOTHER?

Business

Here is something to think about for the new year. You are the managing director of a firm, with regular customers who usually place orders worth hundreds or thousands of pounds, but who sometimes make small purchases as well. Your finance director tells you that the cost of raising an invoice is £25. Do you:

a) Grab every penny?
b) Tell sales to send out orders less than £25 gratis to regular customers?

Being a sensible person you know that option b not only saves you money, but raises the esteem of your company with customers. It's a no-brainer.

Healthcare

Now try the same thing for changing behaviour of health professionals. You are chief executive of a health authority (or whatever they are called nowadays) with particular concern about GP prescribing habits. Let's assume that you have the evidence and good health economic analysis to show that A is not only the same as B, but costs less. Do you:

a) Arrange for hordes of thrusting young pharmacists to tell your GPs that they are wrong?
b) Ask about how savings relate to the cost of making those savings?

The no-brainer answer is b, but the one most often chosen is a. When, to save a few pennies, GPs are pressured to prescribe codeine to patients and tell patients to buy their own paracetamol, rather than prescribing paracetamol and codeine together, we need a bit of common sense. That comes from a group of British health economists [1].

Healthy thinking

In this paper [1] the health economists make the point that the cost of a health intervention is the sum of the cost or saving of the intervention itself, plus the cost of implementing the new policy. Both costs have to be taken into account in calculating cost effectiveness.

Table 1: Cost and effectiveness

ACE inhibitor	
Treatment cost effectivemess (£/QALY)	1437
Implementation cost effectiveness (£/QALY)	297
Total	1734
SSRI	
Cost saving per episode in change to older treatment (£)	50
Cost per patient of change (£)	55
Loss per conversion (£)	5

Their examples are increasing use of ACE inhibitors in patients with heart failure, and the use of older antidepressants instead of SSRIs. The methods are straightforward, and based on costs and benefits seen in a randomised trial of pharmacy advice outreach applied to a health authority.

The results (Table 1) show that for ACE inhibitors, the cost per quality adjusted life year gained is increased, but is still very reasonable. Changing GPs from using SSRIs to older antidepressants actually costs more than it saves.

Comment

This is welcome stuff. There's an old rule that 80% of your savings will come from 20% of accounts. What this paper does is to remind us that there is a cost to changing behaviour. If the effort is not effective, or if it concentrates on trivial targets, then it is a waste of space. Intelligent targeting would be the equivalent of not bothering to raise an invoice for anything costing less than £25. Prescribing advice (often seen as prescribing command) would benefit from this kind of thinking.

References:
1 J Mason et al. When is it cost-effective to change the behaviour of health professionals? JAMA 2001 286: 2988-2992.

INCIDENCE OF PERTHES' DISEASE

Someone once told ***Bandolier*** that 6% of us have a rare disease. Those rare diseases are many and varied, so many and varied, in fact, that the majority of us will rarely or never have heard of more than a few of them. So when ***Bandolier*** was asked to write about the epidemiology, diagnosis and treatment of Perthes' disease the first question was "What's Perthes' disease?".

Perthes' disease is a developmental problem of the hip joint, usually unilateral, affecting younger children, and recognised by history, examination and radiological changes. Literature searches provided no systematic evidence about diagnosis or treatment, but did show up some interesting population studies in the UK that can help thinking about how often it can be expected and temporal trends.

Northern Ireland

Northern Ireland has a centralised orthopaedic service where four hospitals serve children. Over a seven-year period 313 children were diagnosed with Perthes' disease. Postcodes allowed spatial analysis by rural or urban area, and by deprivation index for each postcode district.

The results showed that the overall annual incidence was 11.6 per 100,000 children aged under 15 years. There were 256 boys and 57 girls (4.5 to 1), with mean age at onset of 5.7 years and 16% of cases bilateral.

There was no relationship between settlement size and incidence, and no relationship between population density and incidence. Incidence was highest in children living in

Figure 1: Northern Ireland - Perthes' disease incidence by deprivation index

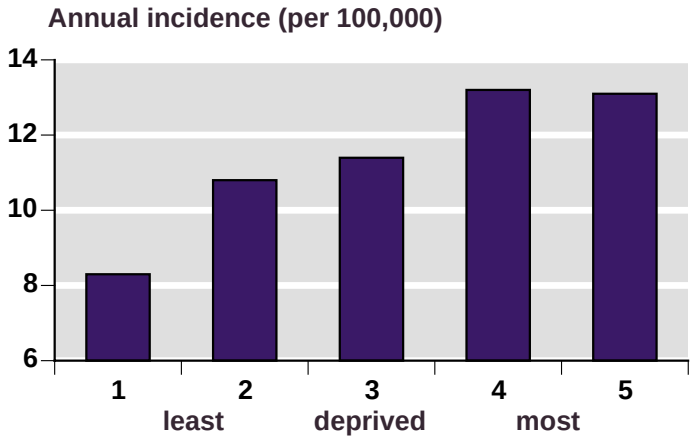

the most deprived areas (Figure 1). This was particularly noted for rural deprivation, where the gradient was steepest (7.1 per 100,000 to 16.1 per 100,000 for increasing deprivation), and for settlements up to 50,000 people. The relationship was not seen in larger settlements.

Liverpool

Liverpool city was the most deprived of 310 English districts in 1998, with the neighbouring districts of Knowsley ninth and Sefton 54th. There is a high incidence of Perthes' disease. Children from all three districts are referred to the Alder Hey Hospital, which has maintained a Perthes' disease register since 1978. The register has been reviewed up to 1999, supplemented by computer searches of other activities in the hospital to ensure no cases were missed.

Parents of affected children were interviewed to determine the district and ward in which the children were born. Denominators for number of children under 15 in wards and districts were determined from census figures. The average ward- and district-specific rates of incidence were calculated for the periods 1976-1981, 1982-1989, and 1990-1995.

Figure 2: Liverpool - Perthes' disease incidence by time

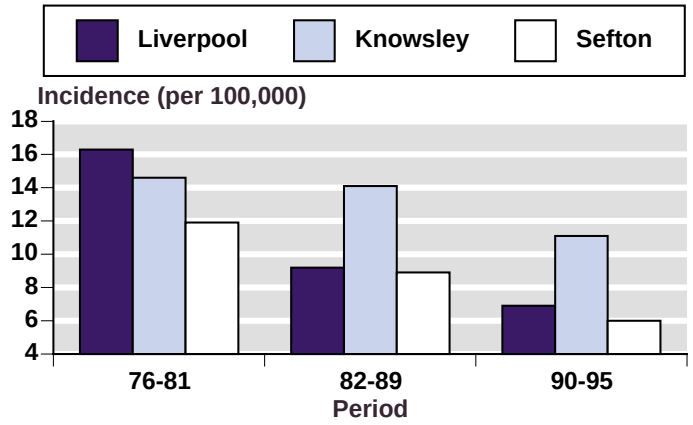

EDITORS

Andrew Moore Henry McQuay
Pain Research
The Churchill, Oxford OX3 7LJ
Editorial office: 01865 226132
Editorial fax: 01865 226978
Email: andrew.moore@pru.ox.ac.uk
Internet: www.ebandolier.com
ISSN 1353-9906

There were 122 children diagnosed with Perthes' disease in Liverpool, 60 in Knowsley, and 38 in Sefton. Incidence rates declined in all three districts (Figure 1). Lower incidence of Perthes' disease was seen in Sefton, the least deprived district (Table 1). Multiple regression analysis was used to summarise various measures of social deprivation and health on ward-specific rates. Incidence increased where deprivation had increased, the prevalence of low birthweight was highest, where free school meals were highest and where wards had a low health index.

Comment

Both studies make general points relating to rare diseases. Obvious is that both ascertained all of the cases in their districts, and went to great lengths to do this and ensure consistency in diagnosis. They went to great lengths to ascertain the proper denominator for their incidence calculations. They went to great lengths to determine accurately factors like deprivation and other indices of deprivation. The result is that we can look on the results with some confidence, despite the numbers of cases in total being under 600.

The message for Perthes' disease is that deprivation is the key. Not urban deprivation, but rural deprivation also, making the results important for Southern counties in the UK which are superficially affluent but where pockets of rural deprivation occur.

Liverpool and Ireland have family links, and there may be genetic components underlying these results. For instance in Liverpool 6% of the population is formed by black or minority groups, but all 122 children diagnosed were white. For a PCO of 100,000 people where about 20% may be under 15 years, high incidence rates like those seen in Liverpool and Northern Ireland would result in two or three children with Perthes' disease every year. Less deprived areas may see only one.

References:

1 WD Kealey et al. Deprivation, urbanisation and Perthes' disease in Northern Ireland. Journal of Bone and Joint Surgery 2000 82-B: 167-171.
2 BM Margetts et al. The incidence and distribution of Legg-Calvé-Perthes' disease in Liverpool, 1982-95. Archives of Diseases in Childhood 2001 84: 351-354.

Table 1: Liverpool - Perthes' disease incidence by district and deprivation index

	Liverpool	Knowsley	Sefton
Perthes' disease per 100,000 children under 15 (1990-1995)	8.7	11.3	4.4
Deprivation index (higher values are more deprived)	10.5	9.4	4.4